arXiv:1608.03628v1 [math.CA] 11 Aug 2016

Time Coupled Diffusion Maps

Nicholas F. Marshall* Matthew J. Hirn†

August 15, 2016

Abstract

We consider a collection of n points in $\mathbb{R}^d$ measured at m times, which are encoded in an $n \times d \times m$ data tensor. Our objective is to define a single embedding of the n points into Euclidean space which summarizes the geometry as described by the data tensor. In the case of a fixed data set, diffusion maps (and related graph Laplacian methods) define such an embedding via the eigenfunctions of a diffusion operator constructed on the data. Given a sequence of m measurements of n points, we construct a corresponding sequence of diffusion operators and study their product. Via this product, we introduce the notion of time coupled diffusion distance and time coupled diffusion maps which have natural geometric and probabilistic interpretations. To frame our method in the context of manifold learning, we model evolving data as samples from an underlying manifold with a time dependent metric, and we describe a connection of our method to the heat equation over a manifold with time dependent metric.

1 Introduction

In many machine learning and signal processing tasks, the observable data is high dimensional, but it lies on a low-dimensional intrinsic manifold. In recent years, several manifold learning methods have emerged which attempt to recover the intrinsic manifold underlying datasets. In particular, graph Laplacian methods have become popular due to their practicality and theoretical guarantees [4, 13, 3, 10, 19, 20, 24, 7].

Current graph Laplacian methods implicitly assume a static intrinsic manifold, or equivalently, that the dynamics underlying the data generation process are stationary. For many applications, this stationary assumption is justified, as datasets often consist of a single snapshot of a system, or are recorded over small time windows. However, in the case where data is accumulated over longer periods of time, accounting for changing dynamics may be advantageous.

*(Corresponding author) Department of Mathematics, Yale University, New Haven, CT 06520, USA (nicholas.marshall@yale.edu)

†Department of Computational Mathematics, Science & Engineering and Department of Mathematics, Michigan State University, East Lansing, MI 48824, USA (mhirn@msu.edu)

Furthermore, if a system is particularly noisy, combining a large number of snapshots over time may help recover structure hidden in noise. These observations raise the following question: how can graph Laplacian methods be extended to account for changing dynamics while maintaining theoretical guarantees? In this paper, we propose modeling data with changing dynamics by assuming there exists an underlying intrinsic manifold with a time-dependent metric.

Moreover, although we assume points on the intrinsic manifold are fixed, their geometry, i.e, dependence structure, is allowed change. We can conceptualize samples from a manifold with a time-dependent metric by considering a corresponding point cloud smoothly moving through $\mathbb{R}^d$ produced by isometrically embedding the manifold over time. The evolution of the metric dictates the movement of points, and vice versa. In practice, datasets conforming to this model are commonly encountered, e.g., an RGB video feed consists of a collection of n pixels which move through $\mathbb{R}^3$.

In Section 2 we describe the proposed method using the diffusion maps framework: a popular graph Laplacian framework which is robust to non-uniform sampling [10]. We remark that diffusion maps is related to other manifold learning methods such as Laplacian eigenmaps and spectral clustering. In fact, if data is uniformly sampled from the underlying manifold, diffusion maps is essentially eigenvalue weighted Laplacian eigenmaps [21]. The central object in the diffusion maps framework is a diffusion operator $\mathbf{P}$ which approximates heat flow on the underlying manifold in a precise sense. In Section 3 we show an analogous result for the case when the manifold is endowed with a time-dependent metric. Numerical results on synthetic data are presented in Section 4. Section 5 investigates the continuous analog of the operator $\mathbf{P}$ and the corresponding time couple diffusion distance. Concluding remarks are given in Section 6.

Recently, other methods to combine or compare multiple measurements have emerged, as the problem is often encountered in practice. In the diffusion maps literature, several notable examples exist. Wang et al. [22] presented a method of combining the output from multiple preprocessing algorithms for image classification. Coifman and Hirn [9] presented a method of extending diffusion maps to allow comparisons across multiple measurements of a system, even when such measurements are of different modalities. Lederman and Talmon [14] as well as Lindenbaum, et al [15] developed a method of finding common variables between two different measurements of a system. Our work extends the current diffusion maps literature by considering evolving dynamics rather than facilitating comparison of fixed systems.

2 Time Coupled Diffusion Maps

In this section, we introduce the notion of time-coupled diffusion maps. Let $\mathcal{M}$ be a compact smooth manifold with a smooth 1-parameter family of Riemannian metrics $g(\tau)$, $\tau \in [0, T]$, $T < \infty$. We will refer to the parameter τ as time throughout. For each time τ, let $\iota_\tau : \mathcal{M} \hookrightarrow \mathbb{R}^d$, $\mathcal{M}_\tau = \iota_\tau(\mathcal{M})$, denote an

isometric embedding of $(\mathcal{M}, g(\tau))$ into d-dimensional Euclidean space, where d is fixed for all $\tau \in [0,T]$.

Let $X = \{x_j\}_{j=1}^n \subseteq \mathcal{M}$ denote a finite collection of n points sampled from $\mathcal{M}$ and let $(\tau_0, \tau_1, \ldots, \tau_m)$ be a uniform partition of $[0,T]$ with $\tau_0 = 0$ and $\tau_m = T$. We will assume our data consists of measurements at $\tau_1, \ldots, \tau_m$ so that we have one measurement for each of the intervals $[\tau_{i-1}, \tau_i)$ for $i = 1, \ldots, m$. Specifically, our data will consist of a sequence of m sets $(X_1, \ldots, X_m)$, where $X_i = \iota_{\tau_i}(X) = \{x_j^{(i)}\}_{j=1}^n$. Such data can be represented as an $n \times d \times m$ tensor corresponding to n points in $\mathbb{R}^d$ measured at m times.

Suppose each X_i is distributed over $\mathcal{M}_{\tau_i}$ according to a density $q_{\tau_i} : \mathcal{M}_{\tau_i} \to \mathbb{R}$. We assume that at time τ_1, the n points $X_1 = \{x_j^{(1)}\}_{j=1}^n$ are sampled independently from $\mathcal{M}_{\tau_1}$ according to the density q_{τ_1}. Since no re-sampling occurs, any changes in the densities q_{τ_i}, for $i = 1, \ldots, m$, result from deformations of q_{τ_1} induced by changes in the Riemannian volume of $\mathcal{M}$, which itself is induced from changes in the Riemannian metric $g(\tau)$ over time. In particular, $x_j^{(i)} = \iota_{\tau_i} \circ \iota_{\tau_1}^{-1}(x_j^{(1)})$ and $q_{\tau_i}(x) = q_{\tau_1}(\iota_{\tau_1} \circ \iota_{\tau_i}^{-1}(x))|\det(D(\iota_{\tau_1} \circ \iota_{\tau_i}^{-1}))(x)|$.

We make two remarks. First, even if the initial sampling density is uniform, the changing geometry may alter the density over time. Hence, the density invariant kernel construction described in [13, 10] plays an essential role in our construction. We will assume the minimum sampling density over time is bounded below and the associated error term enters our error analysis as a constant. We refer the reader to a recent paper by Berry and Harlim [6] for a detailed error analysis of variable density kernel constructions. Second, since we have assumed the initial samples X_1 are i.i.d., then for fixed i, each set of samples X_i are also i.i.d since X_i is the continuous, hence measurable, function $\iota_{\tau_i} \circ \iota_{\tau_1}^{-1}$ of the i.i.d. variables X_1.

For each data set X_i, we construct a diffusion operator $\mathbf{P}_{\varepsilon,i}$ following the diffusion maps framework [10]. For completeness, we include the details of the construction of $\mathbf{P}_{\varepsilon,i}$ in the following. First, for each time index i, we define a Gaussian kernel $\mathbf{K}_{\varepsilon,i}$ on X using the measurements X_i,

$$\mathbf{K}_{\varepsilon,i}(x_j, x_k) = \exp\left(-\frac{\|x_j^{(i)} - x_k^{(i)}\|_2^2}{4\varepsilon}\right), \quad x_j, x_k \in X \text{ and } x_j^{(i)}, x_k^{(i)} \in X_i, \quad (1)$$

where $\|\cdot\|_2$ denotes the Euclidean norm in the ambient space $\mathbb{R}^d$. By summing over the second variable of the kernel $\mathbf{K}_{\varepsilon,i}$ we approximate the density q_{τ_i} by

$$\mathbf{q}_{\varepsilon,i}(x_j) = \sum_{k=1}^n \mathbf{K}_{\varepsilon,i}(x_j, x_k), \quad x_j \in X. \quad (2)$$

Using $\mathbf{q}_{\varepsilon,i}$ we will normalize the kernel $\mathbf{K}_{\varepsilon,i}$ as follows

$$\widetilde{\mathbf{K}}_{\varepsilon,i}(x_j, x_k) = \frac{\mathbf{K}_{\varepsilon,i}(x_j, x_k)}{\mathbf{q}_{\varepsilon,i}(x_j)\mathbf{q}_{\varepsilon,i}(x_k)}, \quad x_j, x_k \in X. \quad (3)$$

Then the corresponding diffusion operator $\mathbf{P}_{\varepsilon,i} : \mathbb{R}^n \to \mathbb{R}^n$ is given by

$$(\mathbf{P}_{\varepsilon,i} f)(x_j) = \sum_{j=1}^{n} \frac{\widetilde{\mathbf{K}}_{\varepsilon,i}(x_j, x_k)}{\sum_{l=1}^{n} \widetilde{\mathbf{K}}_{\varepsilon,i}(x_j, x_l)} f(x_j). \tag{4}$$

Each $\mathbf{P}_{\varepsilon,i}$ is a matrix which is a diffusion operator when applied to column vectors on the left, and a Markov operator when applied row vectors on the right. For $1 \leq t \leq m$, define:

$$\mathbf{P}_{\varepsilon}^{(t)} = \mathbf{P}_{\varepsilon,t} \mathbf{P}_{\varepsilon,t-1} \cdots \mathbf{P}_{\varepsilon,1}. \tag{5}$$

Like its constituent components, the matrix $\mathbf{P}_{\varepsilon}^{(t)}$ is a diffusion operator when acting on column vectors from the left and a Markov operator when acting on row vectors from the right. As a Markov operator, $\mathbf{P}_{\varepsilon}^{(t)}$ acts backwards in time, in the sense that $\mathbf{P}_{\varepsilon}^{(t)}$ acts by first applying $\mathbf{P}_{\varepsilon,t}$, second $\mathbf{P}_{\varepsilon,t-1}$ and so on. On the other hand, as a diffusion operator $\mathbf{P}_{\varepsilon}^{(t)}$ acts forward in time, first applying $\mathbf{P}_{\varepsilon,1}$, second applying $\mathbf{P}_{\varepsilon,2}$ and so on. We have chosen the ordering convention in equation (5) to favor the geometric interpretation. We will define a distance between points $x_j, x_k \in X$ via the posterior distributions of $\boldsymbol{\delta}_j^T$ and $\boldsymbol{\delta}_k^T$ under the Markov operator $\mathbf{P}_{\varepsilon}^{(t)}$, where $\boldsymbol{\delta}_j$ is a Dirac distribution centered at x_j, i.e., $\boldsymbol{\delta}_j(x_j) = 1$ and $\boldsymbol{\delta}_j = 0$ elsewhere. More specifically, following [10] we will define a diffusion based distance as the L^2 distance between these posterior distributions weighted by the reciprocal of the stationary distribution of the Markov chain. That is, we define $D_{\varepsilon}^{(t)}$ by

$$D_{\varepsilon}^{(t)}(x_j, x_k) = \|\boldsymbol{\delta}_j^T \mathbf{P}_{\varepsilon}^{(t)} - \boldsymbol{\delta}_k^T \mathbf{P}_{\varepsilon}^{(t)}\|_{L^2(1/\boldsymbol{\pi}_{(t)})}, \tag{6}$$

where $\boldsymbol{\pi}_{(t)}$ is the stationary distribution of $\mathbf{P}_{\varepsilon}^{(t)}$, i.e., $\boldsymbol{\pi}_{(t)}^T \mathbf{P}_{\varepsilon}^{(t)} = \boldsymbol{\pi}_{(t)}^T$, and $\|\cdot\|_{L^2(1/\boldsymbol{\pi}_{(t)})}$ is the weighted L^2 norm with weight function $1/\boldsymbol{\pi}_{(t)}(x_j)$. We refer to $D_{\varepsilon}^{(t)}$ as the *time coupled diffusion distance*, since it extends the diffusion distance in [10] by coupling the evolution time of the data $(X_1, \ldots, X_t)$ with that of the diffusion process governed by $(\mathbf{P}_{\varepsilon,1}, \ldots, \mathbf{P}_{\varepsilon,t})$. Our next objective is to construct a *time coupled diffusion map*, that is, an embedding of X into Euclidean space which preserves the time coupled diffusion distance. For notational brevity, we will suppress the dependence on ε in the following.

Recall, the diffusion map for a static manifold $(\mathcal{M}, g)$ is defined in terms of the eigenvectors and eigenvalues of the transition matrix $\mathbf{P}$, which is constructed in the same manner as (4). The inhomogeneous transition operator $\mathbf{P}^{(t)}$ is also row stochastic, but unlike $\mathbf{P}$, does not necessarily have a complete basis of real eigenvectors. Instead, we define the operator $\mathbf{A}^{(t)}$ by

$$\mathbf{A}^{(t)} = \boldsymbol{\Pi}_{(t)}^{1/2} \mathbf{P}^{(t)} \boldsymbol{\Pi}_{(t)}^{-1/2}, \tag{7}$$

where $\boldsymbol{\Pi}_{(t)}$ denotes the matrix with $\boldsymbol{\pi}_{(t)}$ along the diagonal and zeros elsewhere. First, observe that $\boldsymbol{\pi}_{(t)}^{1/2}$ is both a right and left eigenvector of $\mathbf{A}^{(t)}$ with

eigenvalue 1. In fact, $\mathbf{A}^{(t)}$ has operator norm one and naturally arises when constructing a diffusion framework starting with a Markov chain, see Appendix A. We compute the singular value decomposition (SVD) of $\mathbf{A}^{(t)}$:

$$\mathbf{A}^{(t)} = \mathbf{U}_{(t)}\mathbf{\Sigma}_{(t)}\mathbf{V}_{(t)}^T, \tag{8}$$

where $\mathbf{U}_{(t)}$ is an orthogonal matrix of left singular vectors, $\mathbf{\Sigma}_{(t)}$ is a diagonal matrix of corresponding singular values, and $\mathbf{V}_{(t)}$ is an orthogonal matrix of right singular vectors. Define

$$\mathbf{\Psi}^{(t)} \triangleq \mathbf{\Pi}_{(t)}^{-1/2}\mathbf{U}_{(t)}\mathbf{\Sigma}_{(t)}. \tag{9}$$

Then we assert that the embedding,

$$x_j \mapsto \boldsymbol{\delta}_j^T\mathbf{\Psi}^{(t)}, \tag{10}$$

preserves the distance $D^{(t)}(x_j, x_k)$. Indeed, by the definition of $\mathbf{A}^{(t)}$ from (7)

$$D^{(t)}(x_j, x_k) = \|\boldsymbol{\delta}_j^T\mathbf{P}^{(t)} - \boldsymbol{\delta}_k^T\mathbf{P}^{(t)}\|_{L^2(1/\boldsymbol{\pi}_{(t)})} = \|\boldsymbol{\delta}_j^T\mathbf{\Pi}_{(t)}^{-1/2}\mathbf{A}^{(t)} - \boldsymbol{\delta}_k^T\mathbf{\Pi}_{(t)}^{-1/2}\mathbf{A}^{(t)}\|_{L^2}.$$

Expanding $\mathbf{\Pi}_{(t)}^{1/2}\mathbf{A}^{(t)}$,

$$\begin{aligned}\|\boldsymbol{\delta}_j^T\mathbf{\Pi}_{(t)}^{-1/2}\mathbf{A}^{(t)} - \boldsymbol{\delta}_k^T\mathbf{\Pi}_{(t)}^{-1/2}\mathbf{A}^{(t)}\|_{L^2} &= \|\boldsymbol{\delta}_j^T\mathbf{\Psi}^{(t)}\mathbf{V}_{(t)}^T - \boldsymbol{\delta}_k^T\mathbf{\Psi}^{(t)}\mathbf{V}_{(t)}^T\|_{L^2}, \\ &= \|\boldsymbol{\delta}_j^T\mathbf{\Psi}^{(t)} - \boldsymbol{\delta}_k^T\mathbf{\Psi}^{(t)}\|_{L^2}.\end{aligned}$$

When $t = 1$, the embedding defined in (10) is equivalent to the definition of the standard diffusion map in [10], see the calculations in Appendix A, which are motivated by [13]. Therefore, the time coupled diffusion map is an extension of the diffusion map. As in [10], the first coordinate of the embedding (10) is always constant. This is because $\boldsymbol{\pi}_{(t)}^{1/2}$ is always the top left singular vector of $\mathbf{A}^{(t)}$. Therefore, in order to produce an embedding of X into $\mathbb{R}^l$ for some $l > 0$, we can compute a rank $l + 1$ singular value decomposition in equation (8) and map

$$x_j \mapsto \boldsymbol{\delta}_j^T\mathbf{\Psi}_l^{(t)},$$

where $\mathbf{\Psi}_l^{(t)}$ denotes the matrix consisting of columns 2 through $l+1$ of the matrix $\mathbf{\Psi}^{(t)}$. Since the singular values of $\mathbf{A}^{(t)}$ are of the form $1 = \sigma_0 > \sigma_1 \geq \cdots \geq \sigma_{n-1}$, the truncated embedding will preserve the diffusion distance up to an error on the order of σ_{l+1}. As in other diffusion based methods, in the case where data lies on a low dimensional manifold, we expect the first few coordinates of the embedding to provide a meaningful summary of the data.

It is instructive at this point to make a comparison with the original diffusion maps of Coifman and Lafon [10]. In that setting, one is given samples X of an isometric embedding of a manifold $\mathcal{M}$ with a *static* metric g. Using the same construction outlined in equations (1), (2), (3), and (4), one obtains a Markov matrix $\mathbf{P}$, in which the time index i is suppressed since the metric is

static. Running this homogeneous Markov chain forward t steps is equivalent to composing $\mathbf{P}$ with itself t times, i.e., $\mathbf{P}^t$. Since $\mathbf{P}$ is, by construction, similar to a symmetric matrix, it can be decomposed into a diagonal matrix of real eigenvalues $\mathbf{\Lambda}$ and a basis of right eigenvectors $\mathbf{\Phi}$. Since the left eigenvectors of $\mathbf{P}$ are orthogonal in $L^2(1/\boldsymbol{\pi})$, the diffusion map

$$x_j \mapsto \boldsymbol{\delta}_j^T \mathbf{\Phi} \mathbf{\Lambda}^t, \tag{11}$$

preserves the diffusion distance $D^t(x_j, x_k) = \|\boldsymbol{\delta}_j^T \mathbf{P}^t - \boldsymbol{\delta}_k^T \mathbf{P}^t\|_{L^2(1/\boldsymbol{\pi})}$, where $\boldsymbol{\pi}$ is the stationary distribution of $\mathbf{P}$. A main result of [10] is that as $n \to \infty$ and $\varepsilon \to 0$, we have $\mathbf{P}_\varepsilon^{t/\varepsilon} \to e^{t\Delta}$, where $e^{t\Delta}$ is the Neumann heat kernel on $\mathcal{M}$. Thus, since the eigenfunctions and eigenvalues of $e^{t\Delta}$ give a complete geometric description of the manifold $\mathcal{M}$ [5], the diffusion maps embedding (11) learns the geometry of the manifold from the samples X.

In the time dependent case, we consider a sequence of data $(X_1, \ldots, X_t)$ and construct a corresponding family of operators $(\mathbf{P}_1, \ldots, \mathbf{P}_t)$. Rather than raising a matrix to a power, we compose the family of operators, obtaining $\mathbf{P}^{(t)}$ defined in equation (5). After defining a distance $D^{(t)}$ based on $\mathbf{P}^{(t)}$ we seek a distance preserving embedding. However, since the product of symmetric matrices is not in general symmetric, there is no reason we should suspect $\mathbf{P}^{(t)}$ to have a basis of real eigenvalues and eigenvectors. Therefore, we construct a diffusion maps framework based on the Markov operator using the singular value decomposition of the operator $\mathbf{A}^{(t)}$ as described in equations (7), (8), (9), and (10) . We expect the operator $\mathbf{P}^{(t)}$ encodes some average sense of affinity across the data, and by extension that the time coupled diffusion map is similarly meaningful. In the following, we will make the connection of $\mathbf{P}^{(t)}$ to the data precise by showing that $\mathbf{P}^{(t)}$ approximates the heat kernel on the underlying manifold $(\mathcal{M}, g(\cdot))$, which will be defined in Section 3. Thus, we conjecture that the time coupled diffusion map aggregates important geometrical information of the manifold $(\mathcal{M}, g(\cdot))$ over the time interval $[0, T]$. Numerical results in Section 4 lend credence to this conjecture. There is, however, no theoretical result that directly links geometrical information of $(\mathcal{M}, g(\cdot))$ over arbitrarily long time scales with its heat kernel. The closest results are contained in [1, 2], in which the heat kernel of $(\mathcal{M}, g(\cdot))$ is used to embed the manifold into a single Hilbert space, so that one can observe the flow of $\mathcal{M}$ over time. This result, however, only holds for short time scales.

Before describing the connection of $\mathbf{P}^{(t)}$ to heat flow, we make a brief computational note. In practice, it is common to construct each diffusion matrix $\mathbf{P}_i$ defined in equation (4) as a sparse matrix, by for example, truncating values which fall below a certain threshold. However, if the product

$$\mathbf{P}^{(t)} = \mathbf{P}_t \mathbf{P}_{t-1} \cdots \mathbf{P}_2 \mathbf{P}_1$$

is computed explicitly, there is no reason to expect sparsity will be maintained. Moreover, since each $\mathbf{P}_i$ is a diffusion operator on an assumed underlying k-dimensional manifold $\mathcal{M}$, if each $\mathbf{P}_i$ initially has l nonzero entires in each row, the number of nonzero entires in each row of the product $\mathbf{P}^{(t)}$ could be on the

order of $(l \cdot t)^k$. Therefore, rather than explicitly constructing $\mathbf{P}^{(t)}$, we can treat $\mathbf{P}^{(t)}$ as an operator that can be applied to vectors in order $n \cdot l \cdot t$ operations. Using this approach, the singular value decomposition in equation (8) can be computed by iterative methods such as Lancroz iteration or subspace iteration [17]. When $n^2 > n \cdot l \cdot t$ this approach may provide significant computational savings, as well as reducing the memory requirements of the computation.

3 The heat kernel for $(\mathcal{M}, g(\cdot))$

We begin by stating the definition and properties of the heat kernel for a manifold with time dependent metric. Let $g(\tau), \tau \in [0,T]$, be a smooth family of metrics on a manifold $\mathcal{M}$. The heat equation for such a manifold is:

$$\frac{\partial u}{\partial t} = \Delta_{g(t)} u, \tag{12}$$

where $u : \mathcal{M} \times [0,T] \to \mathbb{R}$. We say that

$$Z : (\mathcal{M} \times [0,T]) \times (\mathcal{M} \times [0,T]) \to \mathbb{R},$$

is the *heat kernel* for $\frac{\partial}{\partial t} - \Delta_{g(t)}$ if the following three conditions hold:

(H_1) $Z(x,\tau;y,\sigma)$ is C^2 in the spacial variables x, y and C^1 in the temporal variables τ, σ,

(H_2) $\left(\frac{\partial}{\partial t} - \Delta_{g(t)}\right) Z(\cdot,\cdot;y,\sigma) = 0$,

(H_3) $\lim_{\tau \searrow \sigma} Z(\cdot,\tau;y,\sigma) = \delta_y$.

Several researchers have studied the heat equation (12) and corresponding heat kernel for a manifold $(\mathcal{M}, g(\cdot))$, in large part due to its relationship with the Ricci flow [12, 8] but also for data analysis [1]. The following properties of the heat kernel have been established (see [12]): first, the heat kernel exists and is the unique positive function satisfying (H_1), (H_2), and (H_3). Additionally, the solution to the initial value problem,

$$\begin{cases} \partial u/\partial t &= \Delta_{g(t)} u, \\ u(x,0) &= f(x), \end{cases} \tag{13}$$

has the integral representation [12, Corollary 2.2]:

$$u(x,\tau) = \int_{\mathcal{M}} Z(x,\tau;y,0) f(y)\, dV(y,0),$$

where $V(\tau)$ is the Riemannian volume of $\mathcal{M}$ at time τ. The heat kernel $Z(x,\tau;y,\sigma)$ also possesses properties analogous to those of the standard heat kernel, such as the semigroup property,

$$Z(x,\tau;y,\sigma) = \int_{\mathcal{M}} Z(x,\tau;\xi,\nu) Z(\xi,\nu;y,\sigma)\, dV(\xi,\sigma), \quad \forall\, x, y \in \mathcal{M},\ \nu \in (\sigma,\tau). \tag{14}$$

Furthermore, its "rows" sum to one, precisely stated as

$$\int_{\mathcal{M}} Z(x,\tau;y,\sigma)\,dV(y,\sigma) = 1, \quad \forall\, x \in \mathcal{M},\, \sigma < \tau \in [0,T]. \tag{15}$$

In the case of a static metric, the associated integral transform of the Neumann heat kernel is known to have the asymptotic approximation [10]:

$$e^{t\Delta} = \lim_{\varepsilon \to 0} (I + \varepsilon\Delta)^{t/\varepsilon} = (1 + \varepsilon\Delta)^{t/\varepsilon} + \mathcal{O}(\varepsilon). \tag{16}$$

For $t \leq T$, we define the associated integral transform of Z as the operator $T_Z^{(t)} : L^2(\mathcal{M}) \to L^2(\mathcal{M})$,

$$T_Z^{(t)} f(x) = \int_{\mathcal{M}} Z(x,t;y,0) f(y)\,dV(y,0). \tag{17}$$

We want to express $T_Z^{(t)}$ in an analogous form to (16); we would expect it to resemble $e^{\int_0^t \Delta_{g(\tau)} d\tau}$. However, as the family $\Delta_{g(\tau)}$ is not in general commutative with respect to composition, it is necessary to consider the *ordered exponential* of $\Delta_{g(\tau)}$ over $[0,t]$, which is the equivalent of the exponential for non-communicative operators, and which is denoted $\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau}$. The ordered exponential can be expressed in terms of the power series:

$$\begin{aligned}
\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau} &= I + \int_0^t \Delta_{g(\tau)}\,d\tau + \int_0^t \Delta_{g(\tau)} \int_0^\tau \Delta_{g(\tau_1)}\,d\tau_1\,d\tau + \cdots \\
&= I + \sum_{i=1}^{\infty} \int_0^t \Delta_{g(\tau_1)} \int_0^{\tau_1} \Delta_{g(\tau_2)} \int_0^{\tau_2} \cdots \int_0^{\tau_{i-2}} \Delta_{g(\tau_{i-1})} \int_0^{\tau_{i-1}} \Delta_{g(\tau_i)}\,d\tau_i\,d\tau_{i-1} \cdots d\tau_2\,d\tau_1,
\end{aligned}$$

where I denotes the identity in the appropriate space. The operator $\mathcal{T}$ can be thought of as enforcing time order in the products for the standard exponential expansion in a sense which is equivalent to the power series definition, cf. [11]. Considering $\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau}$ as a formal power series, it is easy to see that the heat equation is satisfied. Let $\{\phi_l\}_{l \geq 0}$ be the eigenfunctions of $\Delta_{g(0)}$, ordered in terms of decreasing eigenvalue, and define $E_K \triangleq \text{Span}\{\phi_l : 0 \leq l \leq K\}$. Then $\Delta_{g(0)}$ is bounded on E_K, and furthermore as $g(\tau)$ is a smooth family of metrics for $\tau \in [0,T]$, we can conclude $\Delta_{g(\tau)}$ is also bounded on E_K. Hence on E_K the power series for $\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau}$ converges and by uniqueness we conclude $\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau}$ is the heat kernel on E_K. Henceforth we will use

$$T_Z^{(t)} = \mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau} \tag{18}$$

as the notation for the Neumann heat kernel. We remark that E_K could also be defined independent of the differential structure of the manifold as K-band-limited functions on $\mathcal{M}$, cf. [16]. We prefer the former definition as it provides consistency with the definitions in [10].

Recall that n is the number of spatial samples of the manifold $\mathcal{M}$, and that m is the number of temporal measurements on the time interval $[0,T]$. We will

assume that the time interval $[0,T]$ is divided into m intervals $[\tau_{i-1},\tau_i)$ each of length ε where $\tau_0 = 0$ and $\tau_m = T$. For simplicity, we assume that our m measurements are taken at $(\tau_1,\ldots,\tau_m)$. Our main result is that in the limit of large data, both spatially and temporally, the transition operator $\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)}$ converges to the heat kernel:

$$\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)} \to \mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau} \text{ as } n \to \infty \text{ and } \varepsilon \to 0.$$

More precisely:

Theorem 3.1. *Suppose the isometric embedding $\mathcal{M}_\tau \subset \mathbb{R}^d$ of a time-dependent manifold $(\mathcal{M}, g(\tau))$ is measured at a common set $X = \{x_j\}_{j=1}^n \subset \mathcal{M}$ of n points at ε spaced units of time over a time interval $[0,T]$, so that, in particular, we have time samples $(\tau_i)_{i=1}^m \subset [0,T]$ with $\tau_i = i \cdot \varepsilon$ and $m = T/\varepsilon$.*

Then, for any sufficiently smooth function $f : \mathcal{M} \to \mathbb{R}$ and $t \leq T$, the heat kernel $\mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau}$ can be approximated by the operator $\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)}$:

$$\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)} f(x_j) = \mathcal{T} e^{\int_0^t \Delta_{g(\tau)} d\tau} f(x_j) + \mathcal{O}\left(\frac{1}{n^{1/2} \varepsilon^{d/4+1/2}}, \varepsilon \right), \quad x_j \in X. \tag{19}$$

We will prove Theorem 3.1 in the following. Recall that E_K is defined to be the span for the first $K+1$ eigenfunctions of the Laplace-Beltrami operator $\Delta_{g(0)}$.

Lemma 3.2. *On E_K, the heat kernel $\mathcal{T} e^{\int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} d\tau}$ admits the following asymptotic expansion:*

$$\mathcal{T} e^{\int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} d\tau} = I + \varepsilon \cdot \Delta_{g(\sigma)} + \mathcal{O}(\varepsilon^2). \tag{20}$$

Proof. By definition

$$\mathcal{T} e^{\int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} d\tau} =$$
$$I + \int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} d\tau + \int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} \int_\sigma^\tau \Delta_{g(\tau_1)} d\tau_1 d\tau + \int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} \int_\sigma^\tau \Delta_{g(\tau_1)} \int_\sigma^{\tau_1} \Delta_{g(\tau_2)} d\tau_2 d\tau_1 d\tau + \cdots$$

The tail of this series starting with the third term is $\mathcal{O}(\varepsilon^2)$. For the second term, we can separate the contribution of $\Delta_{g(\sigma)}$ from the integral yielding:

$$\mathcal{T} e^{\int_\sigma^{\sigma+\varepsilon} \Delta_{g(\tau)} d\tau} = I + \varepsilon \cdot \Delta_{g(\sigma)} + \int_\sigma^{\sigma+\varepsilon} (\Delta_{g(\tau)} - \Delta_{g(\sigma)}) d\tau + \mathcal{O}(\varepsilon^2).$$

Now we can bound the integral by using the smoothness of $\Delta_{g(\tau)}$, which results from the smoothness of $g(\tau)$. By the smoothness of $\Delta_{g(\tau)}$ with respect to time and the fact that it is bounded on E_K,

$$\left\| \Delta_{g(\tau)} - \Delta_{g(\tau')} \right\| \leq C|\tau - \tau'|, \quad C \in \mathbb{R}^+.$$

Therefore,

$$\left| \int_\sigma^{\sigma+\varepsilon} (\Delta_{g(\tau)} - \Delta_{g(\sigma)}) d\tau \right| \leq \varepsilon \cdot \sup_{\tau \in [\sigma, \sigma+\varepsilon]} \left\| \Delta_{g(\tau)} - \Delta_{g(\sigma)} \right\| \leq C\varepsilon^2,$$

and hence we have:

$$\mathcal{T}e^{\int_\sigma^{\sigma+\varepsilon}\Delta_{g(\tau)}d\tau} = I + \varepsilon\cdot\Delta_{g(\sigma)} + \mathcal{O}(\varepsilon^2),$$

as desired. □

By using the short time expansion from (20), we can develop a long time asymptotic expansion of the heat kernel. In the following, we will write $\Delta_i = \Delta_{g(\tau_i)}$ to simplify notation.

Lemma 3.3. *Let $(\tau_0, \tau_1, \ldots, \tau_m)$ be a uniform partition of $[0,T]$ with step size $\varepsilon > 0$, so that $\tau_i = \varepsilon\cdot i$ and $\tau_m = m\cdot\varepsilon = T$. For $0 < t \le T$, define $\ell \triangleq \min_i\{i = 1,\ldots,m \mid \tau_i \ge t\}$. Then on E_K the heat kernel $\mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau}$ has the expansion*

$$\mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau} = (I+\varepsilon\Delta_\ell)(I+\varepsilon\Delta_{\ell-1})\cdots(I+\varepsilon\Delta_1) + \mathcal{O}(\varepsilon). \tag{21}$$

Proof. First suppose $t = \tau_\ell$. Using the semigroup property (14) of the heat kernel, we have:

$$\mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau} = \mathcal{T}e^{\int_{\tau_{\ell-1}}^{\tau_\ell}\Delta_{g(\tau)}d\tau}\cdot\mathcal{T}e^{\int_{\tau_{\ell-2}}^{\tau_{\ell-1}}\Delta_{g(\tau)}d\tau}\cdots\mathcal{T}e^{\int_{\tau_0}^{\tau_1}\Delta_{g(\tau)}d\tau}. \tag{22}$$

Combining the expansion (22) and Lemma 3.2, yields the following:

$$\begin{aligned}\mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau} &= \left(I+\varepsilon\Delta_\ell+\mathcal{O}(\varepsilon^2)\right)\left(I+\varepsilon\Delta_{\ell-1}+\mathcal{O}(\varepsilon^2)\right)\cdots\left(I-\varepsilon\Delta_1+\mathcal{O}(\varepsilon^2)\right),\\ &= (I+\varepsilon\Delta_\ell)(I+\varepsilon\Delta_{\ell-1})\cdots(I+\varepsilon\Delta_1)+\Xi.\end{aligned}$$

We now analyze the error term Ξ. First note that for small enough ε we have:

$$\|I+\varepsilon\Delta_i\| \le 1, \quad \text{for all } i = 1,\ldots,m.$$

Furthermore, since $[0,T]$ is fixed, the number of temporal measurements m is $m = T/\varepsilon = \mathcal{O}(1/\varepsilon)$. Putting together these two facts, one obtains:

$$\Xi = \sum_{k=1}^{\ell}\binom{\ell}{k}\mathcal{O}(\varepsilon^{2k}) \le \sum_{k=1}^{\ell}\frac{\ell^k}{k!}\mathcal{O}(\varepsilon^{2k}) = \sum_{k=1}^{\ell}\frac{1}{k!}\mathcal{O}(\varepsilon^k) = \mathcal{O}(\varepsilon),$$

which proves (21) for $t = \tau_\ell$.

Now let $\tau_\ell > t$. The previous argument shows:

$$\begin{aligned}\mathcal{T}e^{\int_0^{\tau_\ell}\Delta_{g(\tau)}d\tau} &= \left(I+\varepsilon\Delta_{g(\tau_\ell)}\right)(I+\varepsilon\Delta_{\ell-1})\cdots(I+\varepsilon\Delta_1)+\mathcal{O}(\varepsilon),\\ \mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau} &= \left(I+\varepsilon\Delta_{g(t)}\right)(I+\varepsilon\Delta_{\ell-1})\cdots(I+\varepsilon\Delta_1)+\mathcal{O}(\varepsilon).\end{aligned}$$

Therefore the difference between the two operators is:

$$\|\mathcal{T}e^{\int_0^{\tau_\ell}\Delta_{g(\tau)}d\tau} - \mathcal{T}e^{\int_0^t\Delta_{g(\tau)}d\tau}\| \le \varepsilon\|\Delta_{g(\tau_\ell)} - \Delta_{g(t)}\| + \mathcal{O}(\varepsilon) \le C\varepsilon^2 + \mathcal{O}(\varepsilon) = \mathcal{O}(\varepsilon),$$

and so (21) holds for any $0 < t \le T$.

□

Proof of Theorem 3.1. Recall that each of our diffusion operators $\mathbf{P}_{\varepsilon,i}$ for a fixed i is defined the same as in the standard diffusion maps framework. Hence, by adapting the convergence result from Singer [18], we have that on E_K:

$$\mathbf{P}_{\varepsilon,i} = I + \varepsilon\Delta_i + \mathcal{O}\left(\frac{1}{n^{1/2}\varepsilon^{d/4-1/2}}, \varepsilon^2\right).$$

Let $\ell = \lceil t/\varepsilon \rceil$ and recall our ℓ-step transition operator $\mathbf{P}_\varepsilon^{(\ell)}$ is defined as $\mathbf{P}_\varepsilon^{(\ell)} \triangleq \mathbf{P}_{\varepsilon,\ell}\mathbf{P}_{\varepsilon,\ell-1}\cdots\mathbf{P}_{\varepsilon,1}$. The proof of Lemma 3.3 demonstrates that the error of the product increases by a factor of ε, hence:

$$\mathbf{P}_\varepsilon^{(\ell)} = (I + \varepsilon\Delta_\ell)(I + \varepsilon\Delta_{\ell-1})\cdots(I + \varepsilon\Delta_1) + \left(\frac{1}{n^{1/2}\varepsilon^{d/4+1/2}}, \varepsilon\right).$$

Therefore using Lemma 3.3 we conclude that

$$\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)} = \mathcal{T}e^{\int_0^t \Delta_{g(\tau)}d\tau} + \left(\frac{1}{n^{1/2}\varepsilon^{d/4+1/2}}, \varepsilon\right),$$

on the set E_K, where K is fixed but arbitrary. The result follows by taking the union and closure $\overline{\bigcup_{K>0} E_K}$, as in [10]. □

Remark 3.4. The connection to heat diffusion can be also considered in terms of infinitesimal generators. For the static case, the product $\mathbf{P}_\epsilon^{t/\epsilon}$ approaches a continuous-time Markov operator P_t, which can be represented in terms of a generator G as $P_t = e^{tG}$ where $G = \Delta$. In the above, we have shown that the $\lceil t/\varepsilon \rceil$-step transition operator for our inhomogeneous Markov chain $\mathbf{P}_\epsilon^{(\lceil t/\varepsilon \rceil)}$ approaches the continuous process $\mathcal{T}e^{\int_0^t \Delta_{g(\tau)}d\tau}$, whose time-dependent infinitesimal generator at time t is $\Delta_{g(t)}$.

4 Numerical Example

Suppose that D consists of $n = 10,000$ points sampled uniformly at random from the unit disc in $\mathbb{R}^2$. Define the maps $h, v : \mathbb{R}^2 \to \mathbb{R}^2$ by

$$h(x,y) = \big(x, y\cdot(1-\cos\pi x)\big), \quad \text{and} \quad v(x,y) = \big(x\cdot(1-\cos\pi y), y\big).$$

We refer the images $h(D)$ and $v(D)$ as horizontal and vertical barbells, respectively. In the following, we will define a deformation of the $n = 10,000$ points from a horizontal to a vertical barbell over $m = 9$ times. At time $i = 1$ we define $X_1 = h(D)$, i.e., the points start in the shape of a horizontal barbell. Then over times $i = 1,\ldots,5$ we move the points linearly in space such that at time $i = 5$ the data $X_5 = D$, i.e., the points are uniformly distributed over the unit disc. Over times $i = 5,\ldots,9$ we again move points linearly in space, but such that at time $i = 9$ the data $X_9 = v(D)$, that is, the points finish in the vertical barbell configuration. The sequence of data sets $(X_1,\ldots,X_9)$ can be stored as an $10,000 \times 2 \times 9$ data tensor corresponding to the $10,000$ points in

$\mathbb{R}^2$ measured over the 9 times. Moreover, by fixing the first coordinate of this tensor, we have a 2×9 matrix corresponding to the trajectory of a single point through $\mathbb{R}^2$ over the 9 times. In Figure 4 we plot the data for times $i = 1, 5, 9$. Note how even though points have uniform sample density at time $i = 5$, the changing area of the manifold creates a nonuniform distribution of points at the other times. Specifically, the density of points is higher towards the middle of each barbell function.

We compare three different methods of embedding the data: a standard diffusion map, an average distance diffusion map, and finally the time coupled diffusion map introduced in this paper.

First, for each fixed data set X_i, we compute the standard diffusion map. The first three coordinates ϕ_1, ϕ_2, ϕ_3 of the diffusion map are plotted for X_1,X_5 and X_9 in Figure 4. The standard diffusion map embeddings for $i = 1$ and $i = 9$ accentuate the two clusters occurring in the data at these times. All points in the figure have been colored using the same color map, which helps distinguish the horizontal and vertical barbell diffusion map embeddings, which are otherwise the same. At time $i = 5$, the coordinates ϕ_1 and ϕ_2 recover the circle, while ϕ_3 is higher frequency oscillation.

Second, we create an embedding based on an average distance function over time. Specifically, we define the average distance function for times $i = 1, \ldots, m$

$$\rho_i(x_j, x_k) = \frac{1}{i} \sum_{l=1}^{i} \|x_j^{(l)} - x_k^{(l)}\|_2 \quad \text{for} \quad j, k = 1, \ldots, n.$$

Using this distance function, we construct a diffusion map via equations (1), (2), (3), and (4) where the distance function ρ_i is used in place of the ℓ^2 norm in equation (1). The first three coordinates of the resulting diffusion map $\zeta_1, \zeta_2, \zeta_3$ are plotted in Figure 4. This average distance diffusion map is equivalent to a standard diffusion map at time $i = 1$, but the effect of the averaging is clearly displayed in the subsequent embeddings. In particular at time $i = 9$, the embedding is kite shaped reflecting the average distance between the points.

Third, for $\{X_i\}_{i=1}^1$, $\{X_i\}_{i=1}^5$, and $\{X_i\}_{i=1}^9$ we computed the time-coupled diffusion map and plot the first three coordinates ψ_1, ψ_2, and ψ_3 in Figure 4. Observe that at time $i = 1$ the time coupled diffusion map is equivalent to the diffusion map. However, already at time $i = 5$, the time coupled diffusion map "remembers" the horizontal barbell and combines the horizontal barbell geometry with that of the circle. Finally at time $i = 9$, observe the time coupled diffusion embedding has separated the data into four lines which meet at the origin. In contrast to the average distance embedding, the time coupled diffusion embedding has maintained a concise description of the data similar to that of the barbell at time $i = 1$. The four lines in the time coupled diffusion maps embedding summarize the four classes of the data. Specifically, at time $i = 1$ the data is separated into two classes– left or right– depending on which side the points reside in the horizontal barbell. Similarly, at time $i = 9$, the data is separated into two different classes – up or down – depending on which side the points reside in the vertical barbell. Therefore, each point can be described as

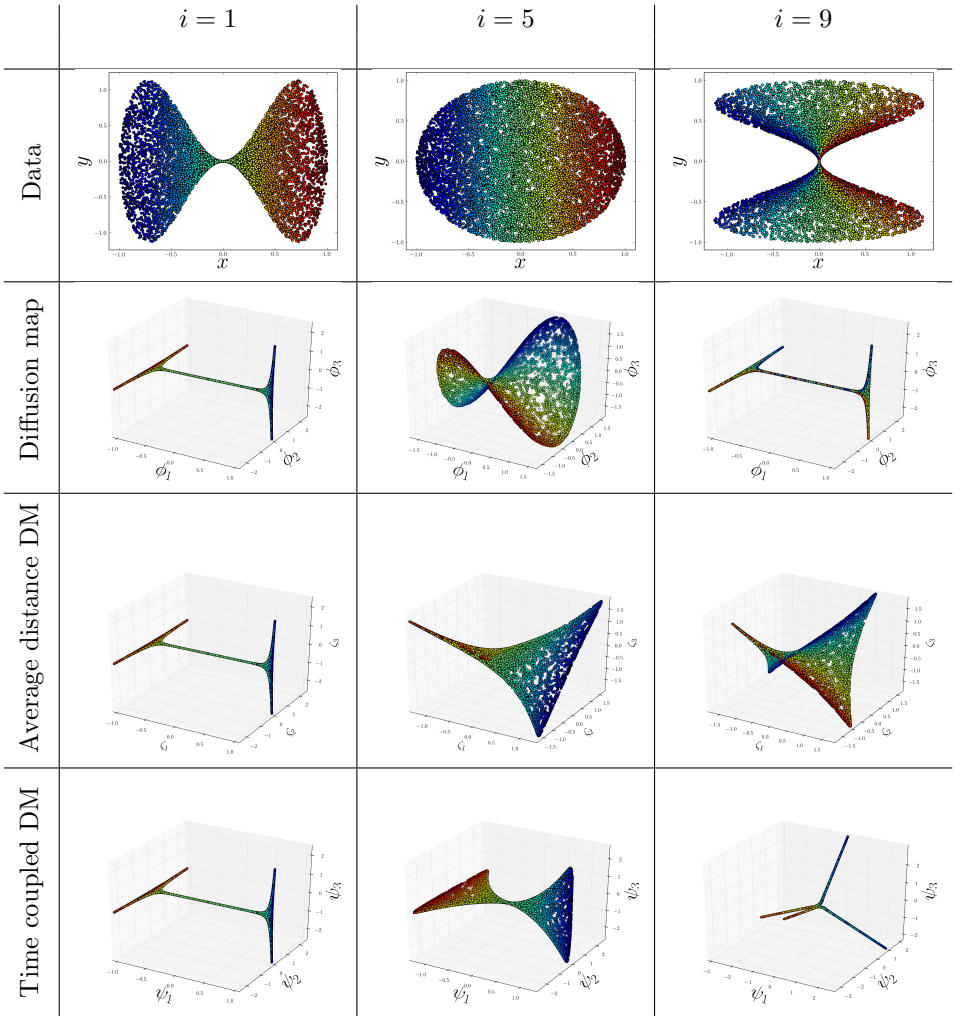

Figure 1: In the first row, the data – consisting of $10,000$ points evolving over 9 time steps, and generated as described Section 4 – is plotted for times $i = 1, 5, 9$. In the second row, the first three coordinates of the diffusion maps for the fixed sets X_1, X_5, and X_9 are plotted. In the third row, the first three coordinates for the distance averaged diffusion map for the sequence of data sets $\{X_i\}_{i=1}^{1}$, $\{X_i\}_{i=1}^{5}$, and $\{X_i\}_{i=1}^{9}$ are plotted. Finally, in the fourth row, the time coupled diffusion map for the sequences of data sets $\{X_i\}_{i=1}^{1}$, $\{X_i\}_{i=1}^{5}$, and $\{X_i\}_{i=1}^{9}$ are plotted. A consistent color map is used across all plots.

either: (right, up), (right, down), (left, up), or (left, down). Each of the four lines in the time coupled diffusion map embedding for time $i = 9$ corresponds to one of these four classes of points. Moreover, since at any given time only two classes of points exist in the data, our example demonstrates how time coupled diffusion maps can summarizing changing geometry over time. To reiterate, no single temporal slice of the data is sufficient to recover the four classes in the data, but the time couple embedding is able to display this information.

5 Extensions to Related Time Dependent Diffusion Operators

Theorem 3.1 approximates the operator $T_Z^{(t)} : L^2(\mathcal{M}) \to L^2(\mathcal{M})$, defined in (17), with the discrete inhomogeneous Markov chain $\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)}$. The inhomogeneous chain $\mathbf{P}_\varepsilon^{(\lceil t/\varepsilon \rceil)}$ is the foundation of the time coupled diffusion distance, and after the normalization in (7) the singular value decomposition can be used to construct the corresponding time coupled diffusion map (10). In Section 5.1 we further investigate the connection between the heat kernel Z and the time coupled diffusion map through the continuous operator $T_Z^{(t)}$.

5.1 Continuous Time Coupled Diffusion Map

Recall the diffusion operator $T_Z^{(t)} f(x) = \int_{\mathcal{M}} Z(x,t;y,0) f(y)\, dV(y,0)$ introduced in Section 3. If f is an initial distribution over $\mathcal{M}$, the operator $T_Z^{(t)}$ diffuses it forward in time over the changing manifold geometry, up to time $\tau = t$. The adjoint of this operator is $(T_Z^{(t)})^* : L^2(\mathcal{M}) \to L^2(\mathcal{M})$,

$$(T_Z^{(t)})^* f(x) = \int_{\mathcal{M}} Z(y,t;x,0) f(y)\, dV(y,t).$$

It follows from (15) that the operator $(T_Z^{(t)})^*$ maps a probability distribution over $(\mathcal{M}, g(t))$ backwards in time to a probability distribution over $(\mathcal{M}, g(0))$. Indeed, if f is a probability density function, then $(T_Z^{(t)})^* f$ is its posterior probability distribution as in Section 2. In the continuous setting, the analysis is performed over the manifold $(\mathcal{M}, g(\cdot))$, and all integration is performed with respect to the Riemannian volume on the manifold similar to the approach in [2, 1, 5]. In the continuous case, the stationary distribution is constant and the normalization in (7) is therefore unnecessary. Moreover, the operator $T_Z^{(t)}$ is bi-stochastic in the sense that the constant function is an eigenfunction of both $T_Z^{(t)}$ and $\left(T_Z^{(t)}\right)^*$. Indeed, in Lemma 3.2 we showed that $T_Z^{(t)}$ can be written as the limit of a product – with an increasing number of terms – of self-adjoint operators, which each have the constant eigenfunction. Therefore, the continuous analog of the discrete diffusion distance (6) is the diffusion distance

$$D^{(t)}(x,y)^2 = \int_{\mathcal{M}} (Z(x,t;w,0) - Z(y,t;w,0))^2\, dV(w,0),$$

or equivalently, $D^{(t)}(x,y) = \|(T_Z^{(t)})^*\delta_x - (T_Z^{(t)})^*\delta_y\|_{L^2(\mathcal{M})}$, where δ_x is the Dirac distribution centered at $x \in \mathcal{M}$ (compare to (6)).

Analogous to the discrete decomposition (8), $T_Z^{(t)}$ is decomposed as:

$$T_Z^{(t)} f(x) = \sum_{k\geq 0} \sigma_k^{(t)} \langle f, \varphi_k^{(t)} \rangle \psi_k^{(t)}(x),$$

where $\varphi_k^{(t)} \in L^2(\mathcal{M})$ is a right singular function of $T_Z^{(t)}$, $\psi_k^{(t)} \in L^2(\mathcal{M})$ is a left singular function of $T_Z^{(t)}$, and $\{\sigma_k^{(t)}\}_{k\geq 0}$ are the corresponding singular values. However, recall that in the discrete case, we must take into account the stationary distribution of the Markov chain using the normalization in (7), since the discrete operator $\mathbf{P}_\varepsilon^{(t)}$ is not in general bi-stochastic. In the continuous setting left singular functions are used to define the time coupled diffusion map (10), written here as:

$$\Psi^{(t)}(x) = \left(\sigma_k^{(t)} \psi_k^{(t)}(x)\right)_{k\geq 1}.$$

where the index k starts at 1 because the 0th left singular function is constant. These functions are eigenfunctions of $T_Z^{(t)}(T_Z^{(t)})^*$. By direct calculation, we can write this operator as:

$$T_Z^{(t)}(T_Z^{(t)})^* f(x) = \int_{\mathcal{M}} K_Z^b(x,t;y,t) f(y)\, dV(y,t), \tag{23}$$

where

$$K_Z^b(x,t;y,\sigma) = \int_{\mathcal{M}} Z(x,t;w,0) Z(y,\sigma;w,0)\, dV(w,0).$$

The kernel K_Z^b measures the similarity between $x \in (\mathcal{M}, g(t))$ and $y \in (\mathcal{M}, g(\sigma))$ by integrating the overlap of the posterior distributions of δ_x and δ_y at time $\tau = 0$. The operator (23) restricts to the time diagonal $\sigma = t$ and fixes t. The left singular functions $\{\psi_k^{(t)}\}_{k\geq 1}$ thus summarize geometrical information of $(\mathcal{M}, g(t))$ by looking backward over the time interval $[0,t]$. By Theorem 3.1, $T_Z^{(t)}$ is approximated by $\mathbf{P}_\varepsilon^{(t)}$, and so these left singular functions are analogous to the left singular vectors computed in (8).

The right singular functions $\{\varphi_k^{(t)}\}_{k\geq 1}$ have an analogous, but forward looking, interpretation. Indeed, they are eigenfunctions of $(T_Z^{(t)})^* T_Z^{(t)}$, which is the integral operator with kernel:

$$K_Z^f(x,y;t) = \int_{\mathcal{M}} Z(w,t;x,0) Z(w,t;y,0)\, dV(w,t).$$

Unlike K_Z^b, the kernel K_Z^f compares $x, y \in (\mathcal{M}, g(0))$ by integrating the overlap of the average distributions that result from diffusing the Dirac distributions δ_x and δ_y forward in time, from time $\tau = 0$ to time $\tau = t$.

Since the time coupled diffusion map is based on the SVD of $T_Z^{(t)}$, we remark that diffusing $\varphi_k^{(t)}$ forward in time results in $\sigma_k^{(t)} \psi_k^{(t)}$, and conversely the

backward propagation of $\psi_k^{(t)}$ leads to $\sigma_k^{(t)}\varphi_k^{(t)}$; specifically $T_Z^{(t)}\varphi_k^{(t)} = \sigma_k^{(t)}\psi_k^{(t)}$ and $(T_Z^{(t)})^*\psi_k^{(t)} = \sigma_k^{(t)}\varphi_k^{(t)}$. If we define,

$$\widetilde{\Phi}^{(t)}(x) = \left(\varphi_k^{(t)}(x)\right)_{k\geq 1}$$

and set $T_Z^{(t)}\widetilde{\Phi}^{(t)}(x) = (T_Z^{(t)}\varphi_k^{(t)}(x))_{k\geq 1}$, then this observation leads to:

$$\begin{aligned} D^{(t)}(x,y) &= \|(T_Z^{(t)})^*\delta_x - (T_Z^{(t)})^*\delta_y\|_{L^2(\mathcal{M})}, \\ &= \|\Psi^{(t)}(x) - \Psi^{(t)}(y)\|_{\ell^2}, \\ &= \|T_Z^{(t)}\widetilde{\Phi}^{(t)}(x) - T_Z^{(t)}\widetilde{\Phi}^{(t)}(y)\|_{\ell^2}. \end{aligned}$$

We see that the difference in posterior distributions of δ_x and δ_y is equivalent to the difference between the diffused distributions of $\widetilde{\Phi}^{(t)}(x)$ and $\widetilde{\Phi}^{(t)}(y)$; these equivalent distances are, in turn, equal to the Euclidean distance between points under the time coupled diffusion map $\Psi^{(t)}$.

6 Conclusion

We have introduced the notion of time coupled diffusion maps as a method of summarizing evolving data via an embedding into Euclidean space. In particular, we have introduced a method of modeling evolving data as samples from a manifold with a time dependent metric. We show that the constructed time inhomogeneous Markov chain approximates heat diffusion over a manifold $(\mathcal{M}, g(\cdot))$ with a smoothly varying family of metrics $g(\tau)$. In the context of manifold learning, these operators and resulting embeddings are related to the heat kernel of $\partial_t u = \Delta_{g(t)}u$, which provides geometric and probabilistic interpretations. Numerical experiments indicate that the map encodes aggregate geometrical information over arbitrarily long time scales, and thus summarizes the geometry of a sequence of datasets in a useful way.

These results open new directions related to diffusion based manifold learning, and in particular the use of inhomogeneous Markov chains and asymmetric diffusion semi-groups to understand data geometry. Early results on single cell data [23] show the usefulness of this type of diffusion process in biology, but further numerical and theoretical study is needed. Understanding precisely the nature of the geometrical information encoded by the time coupled diffusion map (extending [1, 2]), would give theoretical insight and could lead to further developments. More immediately, the results contained here provide a theoretical foundation for understanding dynamic data that can be modeled as a time varying manifold.

7 Acknowledgements

N.M. was a participant in the 2013 Research Experience for Undergraduates (REU) at Cornell University under the supervision of M.H. During the REU

program both were supported by the National Science Foundation grant number NSF-1156350. This paper is the result of work started during the REU. M.H. was supported by the European Research Council (ERC) grant InvariantClass 320959 during its writing. He is currently supported as an Alfred P. Sloan Research Fellow and by startup funds from Michigan State University.

Both authors would like to thank Ronald Coifman for numerous insightful conversations.

References

[1] Hiba Abdallah. *Processus de Diffusion sur un Flot de Variétés Riemannniennes*. PhD thesis, L'Universite de Grenoble, 2010.

[2] Hiba Abdallah. Embedding Riemannian manifolds via their eigenfunctions and their heat kernel. *Bulletin of the Korean Mathematical Society*, 49(5):939–947, 2012.

[3] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, December 2008.

[4] Mikhail Belkin. *Problems of Learning on Manifolds*. PhD thesis, University of Chicago, 2003.

[5] P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994.

[6] Tyrus Berry and John Harlim. Variable bandwidth diffusion kernels. *Applied and Computational Harmonic Analysis*, 1(0):–, 2015.

[7] Tyrus Berry and Timothy Sauer. Local kernels and the geometric structure of data. *Applied and Computational Harmonic Analysis*, page In press, 2015.

[8] Bennett Chow, Sun-Chin Chu, David Glickenstein, Christine Guenther, James Isenber, Tom Ivey, Dan Knopf, Peng Lu, Feng Luo, and Lei Ni. *The Ricci Flow: Techniques and Applicatoins Part III: Geometric-Analytic Aspects*, volume 163. AMS, 2008.

[9] Ronald R. Coifman and Matthew J. Hirn. Diffusion maps for changing data. *Applied and Computational Harmonic Analysis*, 36(1):79 – 107, 2014.

[10] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5 – 30, 2006.

[11] P.-L. Giscard, K. Lui, S. J. Thwaite, and D. Jaksch. An exact formulation of the time-ordered exponential using path-sums. *ArXiv e-prints*, October 2014.

[12] Christine M. Guenther. The fundamental solution on manifolds with time-dependent metrics. *The Journal of Geometric Analysis*, 12(3):425–436, 2002.

[13] Stéphane Lafon. *Diffusion Maps and Geometric Harmonics*. PhD thesis, Yale University, 2004.

[14] Roy R. Lederman and Ronen Talmon. Common manifold learning using alternating-diffusion. Technical report, Yale, 2014.

[15] Ofir Lindenbaum, Arie Yeredor, Moshe Salhov, and Amir Averbuch. Multiview diffusion maps. arXiv:1508.05550, 2015.

[16] Saman Mousazadeh and Israel Cohen. Out-of-sample extension of band-limited functions on homogeneous manifolds using diffusion maps. *Signal Processing*, 108(0):521 – 529, 2015.

[17] Heinz Rutishauser. Computational aspects of f. l. bauer's simultaneous iteration method. *Numer. Math.*, 13(1):4–13, March 1969.

[18] A. Singer. From graph to manifold laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128 – 134, 2006. Special Issue: Diffusion Maps and Wavelets.

[19] Amit Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, July 2006.

[20] Amit Singer and Hau-tieng Wu. Vector diffusion maps and the connection laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144, 2012.

[21] R. Talmon, I. Cohen, S. Gannot, and R. Coifman. Diffusion maps for signal processing: A deeper look at manifold-learning techniques based on kernels and graphs. *Signal Processing Magazine, IEEE*, 30(4):75–86, July 2013.

[22] Bo Wang, Jiayan Jiang, Wei Wang, Zhi-Hua Zhou, and Zhuowen Tu. Unsupervised metric fusion by cross diffusion. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2997–3004, June 2012.

[23] Tobias Welp, Guy Wolf, Matthew Hirn, and Smita Krishnaswamy. A diffusion-based condensation process for multiscale analysis of single cell data. In *ICML Workshop Computational Biology*, pages 1–5, New York, NY, 2016. 5 pages.

[24] Guy Wolf and Amir Averbuch. Linear-projection diffusion on smooth euclidean submanifolds. *Applied and Computational Harmonic Analysis*, 34(1):1–14, January 2013.

A Diffusion geometry from Markov kernels

We show that the operator $\mathbf{A}^{(t)}$, defined in Section 2, line (7), has operator norm one and that $\boldsymbol{\pi}_{(t)}^{1/2}$, the square root of its stationary distribution, is an eigenvector with eigenvalue one. We do so by following [13] and developing a diffusion geometry framework starting from a Markov kernel.

Suppose that $(\mathcal{X}, dx)$ is a measure space equipped with a Markov kernel $p(x, y)$. Moreover, we assume the kernel $p(x, y)$ nonnegative and has the conservation property

$$\int_{\mathcal{X}} p(x, y)\, dy = 1 \quad \text{for all } x \in \mathcal{X}.$$

The kernel $p(x, y)$ induces a Markov operator P^* from $L^2(\mathcal{X})$ to itself defined by

$$(P^* f)(x) = \int_{\mathcal{X}} p(y, x) f(y)\, dy \quad \text{for all } x \in \mathcal{X}.$$

The notation P^* has been used to bring attention to the fact that P^* is the adjoint of the operator P defined as:

$$(Pf)(x) = \int_{\mathcal{X}} p(x, y) f(y)\, dy.$$

We assume that P^* has a unique strictly positive stationary distribution $v^2(x)$, that is to say,

$$(P^* v^2)(x) = \int_{\mathcal{X}} p(y, x) v^2(y)\, dy = v^2(x) \quad \text{for all } x \in \mathcal{X}.$$

For example, if $\mathcal{X}$ is a finite set with the counting measure (as in Section 2), it would suffice to assume that the Markov chain is irreducible and positive recurrent. By the function $v(x)$, we denote the positive pointwise square root of $v^2(x)$. We define an operator $A : L^2(\mathcal{X}) \to L^2(\mathcal{X})$ by

$$(Af)(x) = \int_{\mathcal{X}} \frac{v(x)p(x, y)}{v(y)} f(y)\, dy = \int_{\mathcal{X}} a(x, y) d(y)\, dy \quad \text{for all } x \in \mathcal{X}. \tag{24}$$

Here the kernel $a(x, y) = v(x)p(x, y)/v(y)$. We will refer to A has an averaging or diffusion operator and show it has several nice properties.

Lemma A.1. *The operator norm $\|A\| = 1$, and the norm is achieved by the function $v(x)$.*

Remark A.2. Taking $A = \mathbf{A}^{(t)}$ establishes the claim from the beginning of the appendix, which was originally stated after line (7) in Section 2.

Proof. First we evaluate $(Av)(x)$ to check the second part of the assertion and establish a lower bound on the operator norm:

$$(Av)(x) = \int_{\mathcal{X}} v(x)p(x, y)\frac{v(y)}{v(y)}\, dy = v(x) \int_{\mathcal{X}} p(x, y)\, dy = v(x),$$

using the Markov property of the kernel. To establish that 1 is an upper bound on the operator norm we will show that $\langle Af, Af\rangle \leq \|f\|^2$. First, observe that by Cauchy-Schwartz,

$$\left|\int_{\mathcal{X}} \frac{v^2(x)p(x,y)f(y)}{v(y)}\,dy\right| \leq \left(\int_{\mathcal{X}} v^2(x)p(x,y)\,dy\right)^{1/2}\left(\int_{\mathcal{X}} \frac{v^2(x)p(x,y)}{v^2(y)}f^2(y)\,dy\right)^{1/2},$$
$$= v(x)\cdot\left(\int_{\mathcal{X}} \frac{v^2(x)p(x,y)}{v^2(y)}f^2(y)\,dy\right)^{1/2}. \tag{25}$$

Now starting with $\langle Af, Af\rangle$ we multiply and divide by $v^2(x)$ yielding:

$$\langle Af, Af\rangle = \int_{\mathcal{X}} \frac{1}{v^2(x)}\int_{\mathcal{X}} \frac{v(x)^2p(x,y)f(y)}{v(y)}\,dy\int_{\mathcal{X}} \frac{v(x)^2p(x,z)f(z)}{v(z)}\,dz\,dx.$$

By applying the inequality (25) we see that,

$$\langle Af, Af\rangle \leq \int_{\mathcal{X}}\left(\int_{\mathcal{X}} \frac{v^2(x)p(x,y)}{v^2(y)}f^2(y)\,dy\right)^{1/2}\left(\int_{\mathcal{X}} \frac{v^2(x)p(x,z)}{v^2(z)}f^2(z)\,dy\right)^{1/2}dx.$$

Applying Cauchy-Schwartz again,

$$\langle Af, Af\rangle \leq \left(\int_{\mathcal{X}}\int_{\mathcal{X}} \frac{v^2(x)p(x,y)}{v^2(y)}f^2(y)\,dy\,dx\right)^{1/2}\left(\int_{\mathcal{X}}\int_{\mathcal{X}} \frac{v^2(x)p(x,z)}{v^2(z)}f^2(z)\,dz\,dx\right)^{1/2}.$$

Finally, using the fact that $v^2(x)$ is the stationary distribution shows that:

$$\langle Af, Af\rangle \leq \left(\int_{\mathcal{X}} f(y)^2\,dy\right)^{1/2}\left(\int_{\mathcal{X}} f(z)^2\,dz\right)^{1/2} = \|f\|^2,$$

as was to be shown. □